

Amendments to the Specification:

Please replace the paragraph at page 2, lines 11-20, with the following amended paragraph, in which the phrase “The Big Book of Border Gateway Protocol RFCs” is underlined in the original application and is not being added in this amendment (note that the title page is page 1):

Internetworks such as the Internet ~~[[are]]~~ currently ~~comprised of~~ comprise Autonomous Systems, which exchange routing information via exterior gateway protocols. Amongst the most important of these protocols is the Border Gateway Protocol, or BGP. BGPv4 constructs a directed graph of the Autonomous Systems, based on the information exchanged between BGP routers. Each Autonomous System is identified by a unique 16 bit AS number, and, by use of the directed graphs, BGP ensures loop-free routing amongst the Autonomous Systems; BGP also enables the exchange of additional routing information between Autonomous Systems. BGP is further described in several RFCs, which are compiled in The Big Book of Border Gateway Protocol RFCs, by Pete Loshin, which is hereby incorporated by reference.

Please replace the paragraph at page 4, lines 7-11, with the following amended paragraph:

In some embodiments, the routing intelligence unit may be a self-contained device controlling a single edge router. In other embodiments, a single routing intelligence unit controls multiple edge routers. Though the collection of routers ~~[[s]]~~ is coupled to one or more Internet Service Provider (ISP) links, the individual routers may be coupled to one or more ISP links, or to no ISP links.

Please replace the paragraph at page 10, lines 15-23, with the following amended paragraph:

A diagram showing the high level mechanics of the decision maker prefix scheduler is shown in Figure 6. As illustrated in Figure 6, two threads essentially drive the operation of the scheduler. The first thread 600 polls the database for changes in terms of per-SPAL performance, load, or coverage, and decides on which prefix updates to insert in a Priority Queue that holds prefix update requests. The second thread 602 takes items out of the queue in a rate-controlled fashion, and converts the corresponding update requests into an appropriate set of NLRIs (Network Layer

Reachability Information) that it sends to the local routers, and an appropriate set of NLRIs that it sends to the back channel for communication to other Routing Intelligence Units.

Please replace the paragraph at page 12, lines 12-21, with the following amended paragraph:

In some embodiments of the invention, in the first pass, an asynchronous thread goes through all groups in the GROUP_SPAL table, checking whether the NEW_DATA bit is set. This bit is set by the measurement listener in case a new measurement from a /32 resulted in an update of delay, jitter, and loss in the database. Delay, jitter, and loss, also denoted as d , v , and p , are used to compute an application-specific score, denoted by m . The scalar $[[m]]$ m is used to rate application-specific performance; MOS stands for "Mean Opinion Score", and represents the synthetic application-specific performance. In embodiments of the invention, MOS may be multiplied by a degradation factor that is a function of link utilization, resulting in m . (That is, the larger the utilization of a given SPAL, the larger the degradation factor, and the lower the resulting m .)

Please replace the paragraph at page 20, lines 8-21, with the following amended paragraph:

At the outset, a winner set of SPALs is re-computed for P ; this set includes SPALs for which the performance is close to maximal. In some embodiments of the invention, invalid SPALs are excluded from the winner set computation. Bids from remote SPALs under the control of coordinated Routing Intelligence Units may, in embodiments, be included in the winner set computation. Since the bids corresponding to such remote routes are filtered through BGP, they are in units which are compatible with iBGP's ~~LocalPref~~ LOCAL_PREF, which in some implementations is limited to 0-255. Therefore, one possible implementation is to multiply m by 255. The converted quantity is referred to as MSLP. For consistency, the m values computed for local SPALs are also converted to ~~local-pref~~ LOCAL_PREF units. The new winner is then determined to be the set of all SPALs for which MSLP is larger than ~~MSLP_{max}~~ MSLP_{max} - winner-set-threshold, where $MSLP_{max}$ represents the maximum MSLP for that prefix across all available SPALs, and winner-set-threshold represents a customer-tunable threshold specified in ~~LocalPref~~ LOCAL_PREF units. The related pseudo-code is shown below.

Please replace the paragraph at page 30, line 24, to page 31, line 10, with the following amended paragraph:

Prior to forming the NLRIs, the database is updated [[as]] to include the new flap parameters and prefix-SPAL information (i.e., the new current SPAL for that prefix). The BGP update sent to an edge router may be filtered out by the policy on the router. However, assuming the update is permissible, it may be made to win in the router's BGP comparison process. One implementation is to have the edge router [[to]] apply a high Weight value to the incoming update. (Weight is a common BGP knob, supported in most major implementations of the protocol, but it is not in the original protocol specification.) This technique constrains the update so that it gains an advantage only on the router or routers to which the update is directly sent; this is desirable if some other routers are not controlled by a device such as the one described here. It is also possible to send the update with normal BGP attributes which make the route attractive, such as a high ~~LocalPref~~ LOCAL_PREF value.

Please replace the paragraph at page 33, lines 3-19, with the following amended paragraph:

Configuring Routing Intelligence Unit as a Perfect Listener is desirable, as it allows the support of private peerings. For example, unless Routing Intelligence Unit is configured as a Perfect Listener, when Routing Intelligence Unit hears about a prefix, it can't assume that coverage exists for that prefix across all SPALs. Considering the scenario described above, a prefix that the Routing Intelligence Units learns about could be covered by any of the three SPALs the router is connected to. For example, assume that only SPAL 1 has coverage for a given prefix P; in case the Routing Intelligence Unit asserts a performance route for that prefix across SPAL 2, there is no guarantee that the traffic pertaining to that prefix will be transited by the Service Provider to which SPAL 2 is connected (which we denote Provider 2). In case Provider 2 actually has a private peering with Provider X that obeys [[to]] some pre-specified contract, Provider X could well monitor the traffic from Provider 2, and filter all packets that do not conform to that contract. In case this contract namely specifies that Provider X will only provide transit to customers residing on Provider X's network, then the traffic pertaining to Prefix P will be dropped. If Routing Intelligence Unit were a Perfect Listener, it would only assert performance routes for prefixes across SPALs that are determined to have coverage for these prefixes. This behavior may be referred to as "extremely polite."

Please replace the paragraph at page 38, lines 4-7, with the following amended paragraph:

3. The custom mode: in this case, the user can specify minimum and maximum token rates (and, optionally, bucket sizes), as well as conditions on when to move from one token rate to another. Using this custom mode, customers can tailor the prefix scheduler to their exact needs.

Please replace the paragraph at page 38, lines 10-19, with the following amended paragraph:

Even though the priority queue is sized in such a way that the delay spent in the queue is minimized, there is still an order of magnitude between the time scale of the BGP world, at which level decisions are taken, and the physical world, in which edge stats and interface stats are measured. That is, even though the queuing delay is comparable to other delays involved in the process of changing a route, prefix performance across a given link or the utilization of a given link can change much more quickly. For example, a 2 second queuing delay could be appropriate in the BGP world, while 2 seconds can be enough for congestion to occur across a given link, or for the link utilization to go from 25% to 75%. For this reason, in some embodiments of the invention, the winner set is re-evaluated at the output of the priority queue.